International Journal of Computational Intelligence and Applications
Vol. 5, No. 2 (2005) 231–249

METAHEURISTIC APPROACHES TO TRAFFIC GROOMING IN WDM OPTICAL NETWORKS

YONG XU*, SANCHO SALCEDO-SANZ† and XIN YAO‡
The Centre of Excellence for Research in Computational Intelligence
and Applications (CERCIA)
School of Computer Science, The University of Birmingham
Edgbaston, Birmingham, B15 2TT, UK
**y.xu@cs.bham.ac.uk*
†sss@cs.bham.ac.uk
‡x.yao@cs.bham.ac.uk

Received 12 November 2004
Revised 3 May 2005

The widespread deployment of WDM optical networks posts lots of new challenges for network designers. Traffic grooming is one of the most common problems. Efficient grooming of traffic can effectively reduce the overall cost of the network. But unfortunately, traffic grooming problems have been shown to be NP-hard. Therefore, new heuristics must be devised to tackle them. Among these approaches, metaheuristics are among the most promising ones. In this paper, we present a thorough and comprehensive survey on various metaheuristic approaches to the grooming of traffic in both static and dynamic patterns in WDM optical networks. Some future challenges and research directions are also discussed in this paper.

Keywords: WDM optical network; traffic grooming; genetic algorithm; tabu search; simulated annealing.

1. Introduction

Wavelength-division multiplexing (WDM) technology will undoubtedly dominate next-generation backbone transport networks. In WDM optical networks, hundreds of wavelength channels can be established in a single fibre to make the best use of its huge bandwidth. In such a network, the set of lightpaths on the optical layer, which is used for transmitting an optical message between its end nodes, defines the virtual topology; their interconnections are implemented via optical cross-connects (OXCs). Each OXC can switch the optical signal coming in on a wavelength of an input fibre to the same wavelength in an output fibre. Such network architecture makes it flexible to various services and protection requirements. The deployment and design of such WDM optical networks are currently subjects of intensive research and developmental efforts.

In such network architectures, the traffic requirement between each pair of nodes is generally low-rated (e.g. CO-3) but the wavelength's capacity is usually high-rated (e.g. OC-48 or OC-192). The discrepancy between these two requires that a number of such low-rate traffic streams be multiplexed onto high-rate wavelengths to make full use of their capacity and, at the same time, minimise costs to the network. On the other hand, multiplexing/demultiplexing a traffic stream to/from a wavelength requires that the wavelength drop at the two end nodes of it, at which an electronic equipment must be used to process such an operation. In SONET/WDM networks, for example, a SONET add/drop multiplexer (ADM) is required to electronically combine such low-rate traffic streams onto a high-rate circuit (wavelength). Furthermore, such electronic equipments are often very expensive. Therefore it is required that each wavelength drops at only those nodes that have low-rate streams congregated onto it in order to reduce the cost. This is realised by the addition of a wavelength ADM (WADM) at each node. With the WADM, each wavelength is allowed to either drop at a node or optically bypass it with no electronic equipment requirement at that node for that wavelength. A question that arises in such a system is how to groom low-rate traffic streams onto each wavelength so that the amount of electronic equipment needed in the network is minimised while as few wavelengths as possible are used. This problem, often referred to as the traffic grooming (TG) problem, has been proven to be NP-complete[1] and is a topic of much interest in recent literature.[1–6]

The topic of TG has always been one of the popular problems in the optical communications field since it was first proposed in 1998. It was the main focus in the recent special issue on WDM-Based Network Architectures in the *IEEE Journal on Selected Areas in Communications*.[7] The first workshop on TG in WDM Networks was held in San José, CA, USA recently.[8] Comprehensive review has been done in this subject.[4,9] But there has not yet been any article in the literature that provides a review on metaheuristic approaches to the grooming of traffic in WDM optical networks. Since metaheuristics are very good general-purpose search and optimisation tools, they have been widely used to tackle such problems. In this paper we will provide a thorough and comprehensive survey on metaheuristics to various grooming problems. Some future challenges and research directions are also discussed in this paper. One purpose of this paper is to provide the researchers and practitioners in the metaheuristics community with a new application area, and those in the traffic grooming community with a new tool to improve the grooming performance.

The rest of this paper is organised as follows. Section 2 gives the basic concept of grooming in WDM optical networks. Sections 3 and 4 give a thorough survey on metaheuristics to the grooming of both static and dynamic traffic patterns. Some future challenges and research directions of traffic grooming are discussed in Sec. 5. Finally, conclusions are drawn in Sec. 6.

2. Traffic Grooming in WDM Optical Networks

Traffic Grooming (TG) is one of the most important traffic-engineering techniques. It is defined as the allocation of sub-wavelength tributaries onto full wavelength channels in order to achieve efficient utilisation of network resources, for example, by minimising cost or blocking probability, or by maximising revenue. This is a complicated, combinatorial optimisation problem which has been shown to be NP-complete even in a single-hub unidirectional ring network with nonuniform traffic requirements between each pair of nodes.[1] It was first introduced in 1998[10–12] and has since aroused much interest in researchers from both academia and industry due to its high academic and commercial values. It is also one of the most attractive topics in journals related to optical networks.

In WDM optical networks, each fibre transmits a number of wavelengths simultaneously in a WDM manner and at the same time each wavelength can carry several traffic streams in a time-division multiplexing (TDM) manner. By using the ADM, low-rate traffic can be multiplexed onto a high-rate circuit. Four OC-3 traffic streams, for instance, can be multiplexed onto an OC-12 circuit, and 16 OC-3's onto an OC-48, etc. The number of low-rate circuits a wavelength can accommodate is referred to as *traffic granularity.* Since high-speed ADMs are very expensive, they usually dominate the total cost of today's backbone networks. Hence, to reduce such costs, TG is often used to avoid the situation where each wavelength drops at too many nodes in the network.

The benefit of TG in WDM optical networks can be seen from the simple example given below. Consider a four-node unidirectional ring. The traffic requirement between each pair of nodes is eight OC-3's and the wavelength capacity is OC-48. Thus each wavelength can carry two node-pairs' amount of traffic. Since there are six node-pairs of traffic on the ring, three wavelengths are needed.

Figure 1(a) shows the traffic assignment with no grooming, where wavelength λ_1 accommodates traffic requests between node-pairs $1\leftrightarrow 2$ and $3\leftrightarrow 4$, λ_2 $2\leftrightarrow 3$ and $1\leftrightarrow 4$, λ_3 $1\leftrightarrow 3$ and $2\leftrightarrow 4$. In this assignment, every wavelength drops at every node so that three ADMs are needed at every node. 12 ADMs are needed on the ring.

Figure 1(b) gives another assignment with proper grooming. In this scenario, nodes 1, 2, and 4 are equipped with a WADM, which allows only wavelengths that carry traffic to/from that node to be added/dropped at it, bypassing all other wavelengths. The assignment of traffic is as follows:

$$\begin{aligned} \lambda_1&: \quad 1\leftrightarrow 2 \text{ and } 1\leftrightarrow 3;\\ \lambda_2&: \quad 2\leftrightarrow 3 \text{ and } 2\leftrightarrow 4;\\ \lambda_3&: \quad 1\leftrightarrow 4 \text{ and } 3\leftrightarrow 4. \end{aligned}$$

There are also three wavelengths, but since each wavelength drops at only three nodes, nine ADMs are needed this time, and three ADMs are saved compared to Fig. 1(a).

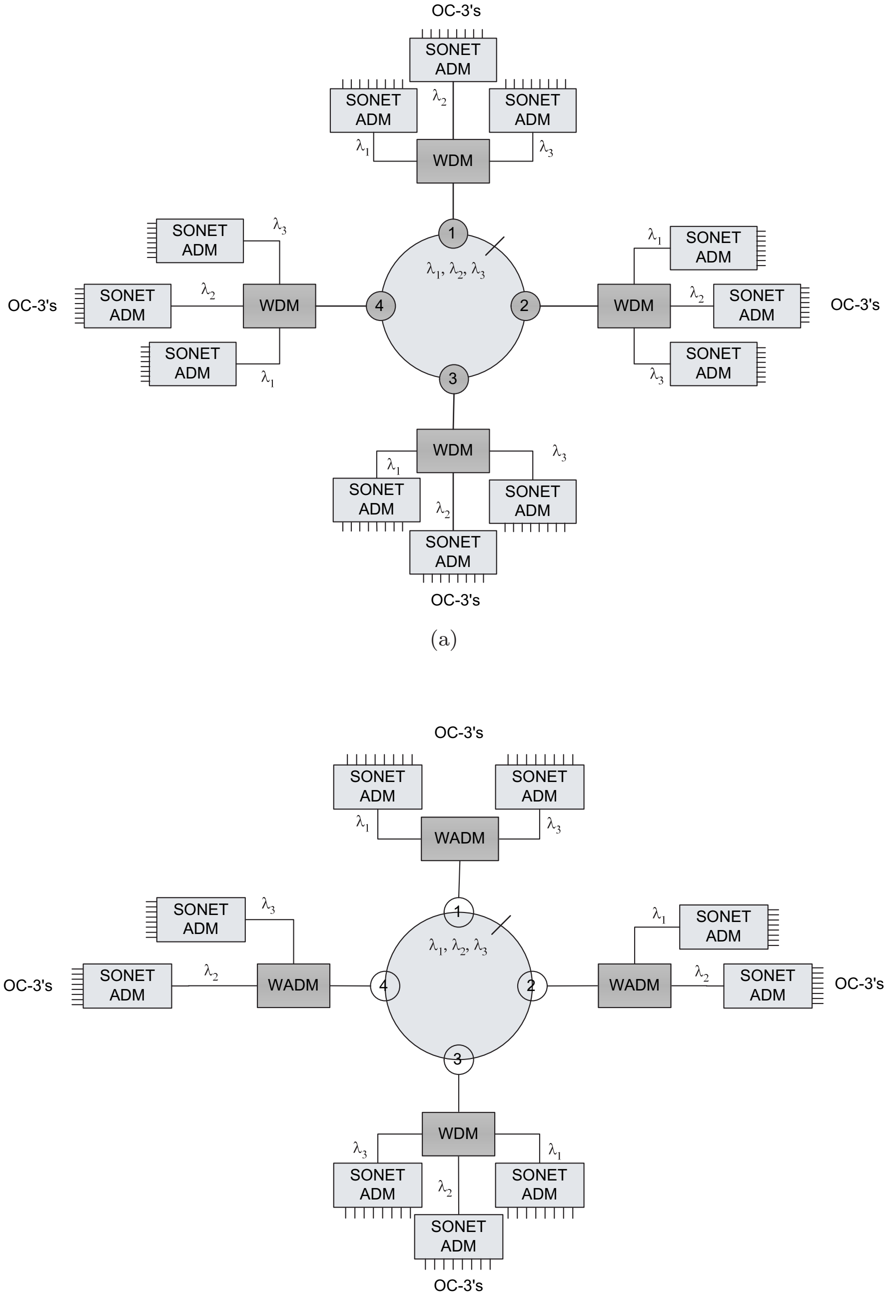

(a)

(b)

Fig. 1. A grooming example on a four-node unidirectional ring with two different scenarios. (a) A possible grooming scenario with 12 ADMs. (b) Another possible grooming scenario with only 9 ADMs.

In these two grooming scenarios, the traffic is carried by the same number of wavelengths. Every wavelength works at their full capacity, but the number of ADMs required is different. This shows that after appropriate grooming, not only can the burden of the electric devices be alleviated, but their numbers can also be reduced. In large-scale networks, the reduction of AMDs due to good grooming will bring about remarkable savings to the costs of constructing them. Therefore, traffic grooming is one of the problems that must be taken into serious account in the design of today's high-speed telecommunications networks.

Depending on whether traffic changes or not, there are basically two grooming categories: static grooming and dynamic grooming. The former is to groom a known and unchanged traffic pattern in order to use the minimum number of ADMs to support it. The latter is to groom a set of changing traffic patterns, either known or unknown, in order to use the minimum number of ADMs to support this changing traffic. It was shown that the traffic grooming problem is NP-hard even for a unidirectional ring with only one egress node and nonuniform traffic. The grooming of arbitrary traffic on a general network topology is even more difficult. Therefore, metaheuristics are good candidates for it. The following will give a thorough survey on metaheuristics to the grooming of static and dynamic traffic.

3. Metaheuristics in the Grooming of Static Traffic

The grooming of static traffic has been one of the most interesting topics in the literature. Although it is more practical to groom dynamic traffic, the grooming of static traffic not only serves as an approximate to the grooming of dynamic traffic, it also provides a basic methodology to it. All the major metaheuristic techniques, i.e. genetic algorithms (GAs), tabu search (TS), and simulated annealing (SA), have been applied to this kind of grooming. We will discuss each of these approaches to this kind of grooming in the following subsections.

3.1. *Genetic algorithms in static traffic grooming*

Due to the power of GAs in combinatorial optimisations, it was the first metaheuristic approach used in the grooming of traffic in WDM optical networks.

In 1999 Tzeng *et al.* proposed a simple GA approach to the grooming of traffic in unidirectional rings.[13] Their objective was to minimise the number of transceivers and wavelengths. They used a chromosome P to denote the state of lightpaths and a sub-chromosome L to indicate the usage of wavelengths in each lightpath. These two chromosomes were arranged in a hierarchical manner to represent the assignment of wavelengths on the ring. A crossover of AND operation was used to combine two sub-chromosomes L_1 and L_2 to produce only one offspring L and a new chromosome P was derived from L. The mutation was to randomly assign a wavelength to carry traffic.

Although this hierarchical chromosome is of interest in that the state of lightpaths and the usage of wavelengths are marked separately, we do not know how to distinguish wavelengths in each lightpath. In addition, when the numbers of lightpaths and wavelengths are large, it could be a very long chromosome. On the other hand, this GA approach has only been applied to a very small problem: an eight-node unidirectional ring with only eight traffic requests. Therefore, we do not know how good this approach is to the grooming of traffic in large-scale networks.

In Ref. 14, Xu *et al.* proposed another GA approach based on the circle-construction model[2] to groom circles onto wavelengths. They first divided every traffic request into basic elements (i.e. the lowest speed elements) then used a heuristic proposed in Ref. 2 to construct C basic circles. Contrary to Ref. 2, in which the authors applied another heuristic to grooming circles onto wavelengths, they employed a GA to grooming circles onto wavelengths. A seed obtained from the greedy circle-grooming algorithm in Ref. 2 was inserted into their initial population and the rest chromosomes were generated by random permutation of the C circles. In decoding they used two methods, next-fit (NF) and best-fit (BF), to assign circles to wavelengths in accordance with their order in the chromosome. NF distributes circles evenly among all wavelengths, but BF tries to fully pack every wavelength except possibly the last one. Two crossovers, partially mapped crossover (PMX) and order crossover (OX), were used to produce offsprings. In mutation, they right-shifted elements of an individual by a random number of positions in a cyclic manner.

Simulations showed that this GA approach achieved better results than the heuristics of Ref. 2, although only grooming results in unidirectional rings were provided, with only a unit connection among all pairs of nodes. Even so, we find that these results are not good enough for large-scale networks. The reasons will be analysed in subsection 5.1. To address these problems, a new GA-based approach has been introduced.

In 2002, Xu *et al.* proposed a new GA-based approach to groom arbitrary nonuniform traffic in SONET/WDM ring networks.[15] Their algorithm operated directly on traffic matrix and groomed traffic in a comprehensive way. They first converted the traffic matrix into a vector then used a random permutation of the traffic items as a chromosome. The position of each item was used to decide the decoding order. They then proposed a first-fit heuristic incorporated with a greedy local improvement to decode each chromosome. The algorithm assigned the first encountered traffic items to each wavelength until one link load on the wavelength exceeded its capacity; the local improvement was then called to assign such a traffic item to it with no additional node added. By assigning a proper fitness value to each chromosome, their approach could be used either to minimise the number of ADMs by using as few wavelengths as possible, or to minimise the number of wavelengths by using as few ADMs as possible. They chose the former as their optimisation objective. The order-mapped crossover (OMX) proposed in Ref. 16 was used to produce offspring and the inversion of a section of some randomly selected genes was used as mutation.

Their algorithm was tested in unidirectional rings with arbitrary traffic requests between each pair of nodes. Simulations showed that their approach were quite good at reducing both the numbers of ADMs and wavelengths. They also found through simulations that the numbers of ADMs and wavelengths could not be minimised simultaneously. One problem with this approach is that they did not consider the splitting of traffic during grooming. Since the amount of traffic requests was not small enough compared to the wavelength's capacity, there were a lot of 'gaps' (unused capacities) left on each wavelength after grooming, which would lead to consuming not only more ADMs but also more wavelengths. Hence, they later on proposed another GA approach incorporated with splitting techniques to further improve the grooming results.[17,18]

3.2. *Grooming with other metaheuristic approaches*

Some other metaheuristics, including TS and SA, are also applied to improve grooming performance in various WDM optical networks with static traffic patterns.

In Ref. 3, Wang *et al.* proposed an SA approach to tackling the static traffic grooming problem in the ring network. They adopted the two-step method used in Ref. 2 to their algorithm. They first used the heuristic in Ref. 2 to construct circles, then randomly combined circles to wavelengths. In the SA process, the authors randomly picked two circles on different wavelengths, swapping part of them or the whole circles as a perturbation. Since the probability for partial swapping was set to be very small, whenever the criterion for partial swapping was satisfied, they checked if there were segments in the two circles that were swappable with no broken connections. If the result was negative, the whole circles were swapped. If the perturbation reduced the ADM usage, it was accepted, otherwise, it would be accepted by probability. The process terminated when some predefined criteria were satisfied.

From the simulations it was found that for most of the time, the SA approach achieved better results than the greedy heuristic in Ref. 2 for uniform traffic with the single-hop case. Sometimes, it even reached the lower bound. But for arbitrary nonuniform traffic, although both approaches achieved good savings in ADM's compared to its upper bound, the SA approach only brought up results similar to the greedy heuristic.

The main problem with this approach is the same for all the others that used the circle-construction method. Since the circles have been fixed in the first step and the minimum number of wavelengths has been used in the second, it is very difficult for this SA approach to obtain optimal results for arbitrary nonuniform traffic patterns. The reasons will be discussed in subsection 5.1.

In Ref. 19, Li *et al.* proposed a similar SA approach to traffic grooming in the ring network. To make their grooming more efficient, the authors introduced traffic-cutting into their algorithm. They first used the method in Ref. 2 and a cut-first heuristic to construct circles, then proposed a hill climbing (HC) and an SA approach to grooming circles onto wavelengths. Following Ref. 20, they defined the operators of merging, splitting, and combining. *Merging* refers to swapping two segments with the same end nodes in two different wavelengths to see if this would bring more end-node sharing. *Combining* refers to finding two segments of unnecessarily split connections and combining them. *Splitting* refers to splitting a connection such that the split parts can fit perfectly into the gaps of other wavelengths. All these operators may bring more end-node sharing and make the assignment more efficient.

HC starts from a current solution, then a set of new solutions is generated by applying the above operators. A new solution is selected randomly from the solutions that are better than the current one. In this way, a good solution can be reached along a certain path. After storing a number of paths, the best solution is returned.

In SA, the same operators as in HC are applied but a different strategy is used to select new solutions. A random solution S is chosen from current solution $S1$ and a new current is made with probability $\min\{1, \exp(-\delta/T)\}$, where $\delta = Cost1 - Cost2$, $Cost1$ is the number of end nodes in solution S, and $Cost2$ is the number of end nodes in solution $S1$. T is the 'temperature' of SA. In this way, local optima can escape. As in HC, SA returns the best solution found over all paths.

Simulations found that both HC and SA approaches produced solutions with nearly 10% fewer end nodes than the cut-first and assign-first (no splitting) approaches alone. The authors also found that HC was not only more stable than SA, but that when the number of nodes on the ring was large, HC could achieve better results than SA. The authors did not specify the reasons. We think the reason may be the same as that stated in Ref. 3, that is, the search space is too large in the nonuniform-traffic case for any SA to find good solutions.

Table 1 summarises various metaheuristics in the grooming of static traffic, from which we are able to know the strengths and weaknesses of each approach.

4. Metaheuristics in the Grooming of Dynamic Traffic

Compared to static traffic grooming, it is more practical and general to groom dynamic traffic patterns since static traffic is either merely a special case or a simplification of dynamic traffic. In the grooming of dynamic traffic, we often use a set of traffic patterns to represent its change. There are mainly two categories of grooming[6]: *strictly nonblocking* and *rearrangeably nonblocking*. If the same traffic request in the set of traffic patterns can be assigned to the same wavelengths without rearranging the assignments among them when traffic changes, it is strictly nonblocking grooming. By this way, all requests in any new traffic pattern can be established without being interrupted. On the contrary, rearrangeably nonblocking

Table 1. Comparison of metaheuristics in the grooming of static traffic.

Algorithm and Method	Strengths	Weaknesses	References
GA, hierarchical chromosome	State of lightpaths and usage of wavelengths are marked separately	Has only been applied to a very small problem	13
GA, circle-construction model	Clear meaning; quite straightforward	Unable to obtain optimal results	14
GA, directly operating on traffic matrix	Has potential to obtain optimal results	Does not consider traffic split	15
SA, circle-construction model	Clear meaning; quite straightforward	Unable to obtain optimal results	3
SA, circle-construction model with traffic-cutting	Clear meaning; quite straightforward	Unable to obtain optimal results	19

grooming is established by assigning each traffic request to the existing wavelengths with possibly rearranging some of the requests among them but the set of nodes at which each wavelength drops is kept unchanged. Due to the possible rearrangement of the requests among different wavelengths, some of the existing traffic may be interrupted when new traffic is established. Nevertheless, this grooming category is usually more efficient in saving ADMs and wavelengths compared to the grooming of strictly nonblocking, since more ADMs can be shared by rearranging some requests in different traffic patterns.

The grooming of dynamic traffic is much more difficult than the grooming of static traffic. Hence, it has not been extensively investigated. As for the metaheuristics, only GA and TS approaches have been applied to this kind of grooming so far. We will give a thorough review of these two metaheuristics to the two dynamic grooming categories described below.

4.1. *Genetic algorithms in dynamic traffic grooming*

GAs are again the first approaches applied to these two grooming categories. In Ref. 21, Xu *et al.* proposed a complicated two-step GA approach to grooming dynamic traffic in the unidirectional ring network in a rearrangeably nonblocking way. Their approach was also based on the circle-construction approach in Ref. 2. They first treated the dynamic traffic patterns at each moment as static ones and followed the heuristic in Ref. 2 to construct circles. They then employed a GA to combine the circles to obtain a static grooming result. Finally, they utilised another GA, by following the ways described in Ref. 5, to match the M virtual topologies, each corresponding to a traffic pattern, to obtain a dynamic grooming result.

At the first step, a chromosome was produced by randomly placing the C circles in a $W \times G$ matrix, where G is the ratio of the wavelength's capacity to the basic traffic rate, and $W = \lceil C/G \rceil$ is the minimum number of wavelengths used on the ring. Each entry of the matrix corresponds to a circle and each row of it denotes the circles that have been assigned to a particular wavelength. Since C/G may not be an integer, there may be some 'empty' circles in the matrix's entries. A modified PMX operator was used as crossover and a random swapping of two circles on two different wavelengths was used as mutation.

After the grooming for every traffic pattern in the first step, they generated another random permutation of a $W \times M$ matrix $\mathbf{T}$ as the chromosome in the second step, where M is the number of traffic patterns. Each element w_{ij} in $\mathbf{T}$ is a binary row vector $w_{ij} = [1/0, 1/0, \ldots, 1/0]_{1\times N}$, representing the placement of ADMs on the ith wavelength of topology j. If there is an ADM at the kth node, $k \in [1, N]$, the corresponding kth element in w_{ij} is 1, otherwise w_{ij} equals 0. The elements on each column were randomly permutated to generate a chromosome. They used a two-point crossover to produce offspring and mutated a chromosome by re-permuting a column.

In simulations, the authors only tested their algorithm with uniform and t-allowable traffic patterns although they claimed that t-allowable traffic was sometimes uneconomical. Another drawback to their approach lies in their three-step grooming process: a heuristic to construct circles; a GA to groom circles onto wavelength; and another GA to match virtual topologies to realise dynamic grooming. Disregarding that it is unlikely for all of these three steps to be optimal, even if optimal results in every step is obtainable, the final result could also be sub-optimal for it is obvious that three optimal results in each step do not necessarily lead to a global optimal solution to the problem. Thus this three-step method is unlikely to achieve optimal grooming results and could be very time-consuming. Furthermore, as was discussed in Ref. 6, this matching method inevitably leads to some idle ADMs for each traffic pattern, i.e. only a part of the ADMs might be used for a certain traffic pattern, which is also a form of wastage in ADM resources. Finally, the use of the minimum number of wavelengths in their grooming does not guarantee that the number of ADMs is minimised at the same time.

Based on the above observations, Xu *et al.* proposed a comprehensive GA approach in Ref. 6 to carry out strictly nonblocking grooming of dynamic traffic in ring networks. The chromosome representation was the same as in Ref. 15, but this time, they employed one chromosome to represent M traffic patterns. In decoding, they proposed a first-fit approach incorporated with a greedy local improvement algorithm to assign traffic items in the chromosome to the current wavelength. If a traffic item could no longer be assigned to the wavelength, a local improvement algorithm was called to assign such a traffic item to it with no additional nodes added. Four different local improvement scenarios were introduced and compared in Ref. 6.

In simulations, they found that it is important to combine GAs with an appropriate heuristic in order to improve grooming performance. They compared their GA approach with the local greedy algorithm only and found that the former significantly outperformed the latter. If a pure GA with no local improvement was used, its performance was worse than the heuristic alone. The benefit of grooming multiple traffic patterns instead of the maximum static traffic was also analysed and demonstrated by simulation. The authors pointed out that the benefit of such grooming is that it can make full use of the dynamics of the changing traffic to make some compensation among the traffic requirements between different pairs of nodes. When traffic changes dynamically, the traffic requirements at their lower level can compensate for those at their higher level. In this way, more ADMs as well as wavelengths can be shared compared to the grooming of maximum traffic. One problem with the authors' approach is that they did not consider the splitting of traffic during grooming which would lead to some 'gaps' on each wavelength. Hence, they proposed another GA approach incorporated with splitting techniques to further improve the grooming results.[17,18]

Reference 17 deals with strictly nonblocking grooming while Ref. 18 deals with the rearrangeably nonblocking one. In these two papers, the authors incorporated three splitting techniques into GA approaches to groom dynamically changing traffic. They analysed the benefit of grooming by splitting, and proposed two splitting methods: traffic-cutting and traffic-dividing. They then combined them to form the third splitting method: synthesised splitting. The first method is to cut traffic into shorter segments so that they are more easily fit into the remaining capacities of the wavelengths. With this method, the original single-hopped traffic becomes multi-hopped and some intermediate nodes are needed to bridge it. Therefore, they suggested not cutting a traffic flow into too many segments. The second method is to divide traffic into smaller parts operating at lower rates so that the lower-rate parts can be assigned to the remaining capacities of the wavelengths. With this method, although the traffic remains single-hopped, the network's virtual topology will become more complex. Hence, the traffic flow should not be divided into too many parts either. The third is a combination of these two splitting methods.

In Ref. 17, the authors' designed a new heuristic to decode the chromosome and combined the GA approach with different splitting techniques to further improve the grooming performance. In this approach, they first assigned a first-encountered traffic item in the chromosome to the current wavelength. The algorithm then examined the remaining traffic items to assign such traffic to the wavelength whose source and termination nodes had already been its dropping nodes. After that, the algorithm tried to assign the next-encountered traffic in the chromosome to the current wavelength and then proceeded with its search again. This process continued until one of the link loads in the current wavelength exceeded its maximum. In this way, when a wavelength dropped at π nodes, the algorithm could assign as many traffic items as possible to it to minimise the numbers of ADMs and wavelengths. They found that this decoding approach was more efficient in reducing the number of ADMs than the one in Ref. 6. If no traffic could be assigned to the current wavelength as a whole, the algorithm would attempt to assign traffic to the current wavelength by one or more splitting methods without adding new ADMs.

The difference between this splitting method and the previous ones in Ref. 20 for dividing and Ref. 19 for cutting is that it cuts and divides traffic only when required. But in Ref. 20 the authors divided all traffic into basic elements and in Ref. 19 some of the basic elements were further cut into segments. This meant that each traffic flow was split into much smaller fragments and each fragment was inclined to be assigned to different wavelengths, which inevitably made a network more complex and increased the control overhead of the system. The signal transmission would be delayed as well. On the contrary, in Refs. 17 and 18, traffic was cut or divided only when there was spare capacity on the wavelength and one of the cuts or divided parts could be assigned to it with no additional nodes added. In this way, the amount of split traffic was minimised and so was the control overhead.

In Ref. 18, the authors incorporated the splitting method with a GA having a hierarchical chromosome structure to tackle the rearrangeably nonblocking grooming problems. In this algorithm, each individual is composed of two hierarchical chromosomes: a master chromosome and a slave chromosome. The use of this hierarchical chromosome structure was based on the fact that in rearrangeably nonblocking grooming each wavelength would drop at the same set of nodes when traffic changed but for different traffic patterns the same traffic demands might be assigned to different wavelengths. Therefore, they used the master chromosome to record each wavelength's dropping nodes and M slave chromosomes to track the assignments of M traffic demands among the wavelengths. Furthermore, the slave chromosomes were only accompanying chromosomes which were produced by a heuristic to rearrange traffic demands among wavelengths and which would then delete some spare ADMs and/or wavelengths after such a rearrangement. The slave chromosomes were updated when a new master chromosome was generated by crossover or mutation.

In simulations, they found that the splitting methods were always more efficient in saving both ADMs and wavelengths than the non-splitting one and that the synthesised one achieved the best results. They also found that the savings were more significant with splitting methods for large-scale networks. They found that when the traffic granularity is large enough, the dividing method is usually better than the cutting method; but when the traffic granularity is small, the cutting method is better than the dividing method in almost all the cases. Comparing their results with those obtained from Ref. 21 they found that this new method is much more efficient in saving ADMs and that the method used in Ref. 21 only led to results worse than no grooming by using the minimum number of wavelengths. Therefore, we have to make a compromise between minimising the number of ADMs and minimising the number of wavelengths when choosing a specific grooming algorithm. The best way is to use a multiobjective optimisation tool to find the Pareto optimal of the problem.

4.2. *Dynamic grooming with tabu search approaches*

Apart from GAs, TS has also been applied to improve grooming performance in ring networks with dynamic traffic patterns.

In Ref. 22, Zhang *et al.* addressed an interesting grooming problem: reconfigurable grooming of dynamic traffic in ring networks. The authors treated dynamic grooming as a series of additional traffic grooming. That is, based on the current wavelength assignment, when traffic changes, a grooming algorithm is used to reconfigure the wavelength assignment according to the new traffic pattern without disrupting the old traffic assignment. Two cases, best-fit and full-fit, for handling reconfigurable SONET over WDM networks were proposed in this paper. *Best-fit* means to place as much new traffic as possible using the available capacity of the current configuration without increasing the number of ADMs, while *full-fit* means to satisfy all new traffic by adding the minimum number of ADMs.

For each case, an integer linear programming model, a greedy and a TS heuristic (TS-1 and TS-2) were given, all of which were based on the two-step circle construction model proposed in Ref. 2.

For the best-fit case, the greedy algorithm starts from grooming the first traffic matrix $\mathbf{T}[i,\ j]$ using the algorithm in Ref. 2. The information regarding whether the entire capacity between two nodes has been occupied, whether there has been an ADM at a node, and at which wavelength a circle has been groomed on, should also be recorded. When a new traffic matrix $\mathbf{T}$'$[i,\ j]$ is given, the difference between the two matrices, $\mathbf{D}[i,j] = \mathbf{T}'[i,j] - \mathbf{T}[i,j]$, is computed. If there are some connections in the old matrix that do not exist in the new one, they are removed. If there are two continuous connections over two circles, they are merged into a bigger one. In grooming new traffic, the traffic in the difference matrix $\mathbf{D}[i,\ j]$ is assigned to the wavelengths from the shortest hop length to the longest hop length with no new ADM added.

The TS heuristic for the best-fit case (TS-1) takes the grooming result from the greedy algorithm as its initial solution. The neighbourhood is defined by swapping part of two circles of the solution in which at least one wavelength should have enough capacity. After such swapping, the assigning of the remaining traffic to the swapped wavelengths is attempted. For all the candidate solutions, the one that is not tabu and which could increase the most traffic is chosen as the new current solution. The search terminates after the pre-specified number of iterations has been reached or if there has been no improvement after a certain number of steps.

A three-phase algorithm was developed for the full-fit case. First, a best-fit algorithm was used (greedy or TS-1) to groom as much traffic as possible to wavelengths with the existing ADMs. Secondly, if there was enough capacity on a wavelength for additional connections, an ADM was placed at a node to groom additional traffic onto it. Finally, another TS heuristic (TS-2) was used to groom the remaining traffic onto new wavelengths. Since the structure of TS-2 is basically the same as TS-1, we neglect it here.

Simulation results have shown that TS-1 could yield better solutions but takes longer than the greedy algorithm for the best-fit case. For the full-fit case, TS-2 yields more competitive results compared to an earlier SA-based method and is more stable for the dynamic case.

The problem with the above algorithms is that either TS-1 or TS-2 is in fact a static offline grooming of changing traffic. The relatively long running time prevents

Table 2. Comparison of metaheuristics in the grooming of dynamic traffic.

Algorithm and Method	Strengths	Weaknesses	References
GA, circle construction model with a two-step GA	None	Unable to obtain optimal results; very time-consuming	21
GA, directly operating on M traffic matrices	Has the potential to obtain optimal results	Does not consider traffic split	6
GA, directly operating on M traffic matrices	Has the potential to obtain optimal results with traffic split	Requires a longer time	17, 18
TS, circle-construction model with best- and full-fit cases	Clear meaning; quite straightforward	Unable to obtain optimal results; difficult to dynamically add new ADM's	22

them from being used in online provisioning of dynamic traffic. Another problem exists, which is also inherent in all the others that used the circle-construction algorithm as their starting point: the number of ADMs and wavelengths cannot be minimised concurrently. An additional problem with the best-fit approach is that the method to dynamically add new ADMs to the network when needed is still not yet known.

Table 2 summarises various metaheuristics in the grooming of dynamic traffic. We can see the strengths and weaknesses of each approach from this table.

5. Problems and Challenges

TG is a newly emerged problem in WDM optical network design. It has drawn much attention from both academia and industry. Efficient grooming of traffic not only reduces the whole network's cost but also increases signal quality by using fewer wavelengths on each fibre. Unfortunately, the traffic grooming problem is a complicated NP-hard problem to which highly efficient heuristics must be utilised. The use of metaheuristic approaches to the grooming problems enables us to achieve high-quality solutions in a reasonable time. However, till now only preliminary work has been done in applying metaheuristic approaches to the grooming problems in which only ring networks have been considered. Therefore, in the future, the following problems must be addressed in this field.

5.1. *Algorithms directly operating on traffic matrices*

We have found in the survey on this topic that most of the approaches that have been used are based on the circle-construction method proposed in Ref. 2. Although it is quite straightforward and reasonable for the circle-construction method to be used in metaheuristic algorithms and despite this method possessing clear meaning, there are still some problems with it. The first is that the numbers of ADMs and wavelengths cannot be minimised concurrently.[1,15] Since this method uses the minimum number of wavelengths as its searching point, the number of ADMs cannot be minimised at the same time for arbitrary nonuniform traffic patterns. The second is that this is a two-step method. Hence, even if the optimal solution can be found at each step, it is still uncertain that the optimal solution for the whole problem can also be found. In addition, with this two-step approach, there are always some idle ADMs.[6] The third problem is that with this method all the connections on the ring have to be divided into elementary parts and different parts in the same connection will probably be assigned to different wavelengths. This will inevitably increase the complexity of the virtual topology and the management overhead of the network. The last is that this method can only be used in ring networks. For more general topologies, e.g. mesh networks, it is difficult to construct circles. This is why most of the metaheuristic approaches that have been used so far deal only with the grooming of traffic in ring networks.

Therefore, it is imperative to design metaheuristic algorithms that operate directly on the traffic matrices and use only a one-step approach to groom traffic onto wavelengths. Not only can these algorithms be used in ring networks, they can also be easily extended to other general topologies. On the other hand, as it has been repeatedly shown that pure (or classical) GAs or other metaheuristic approaches often perform poorly in solving sophisticated combinatorial optimisation problems, an efficient heuristic decoding method must also be imbedded in these algorithms to improve their grooming performance. If splitting techniques have to be included, it is ideal that traffic is split only when needed and that one of the split parts must be assignable to the current wavelength.

5.2. *Grooming in mesh and other topology networks*

Although it is reasonable to groom traffic in ring networks due to its widespread use in today's infrastructural networks, the mesh topology is still more practical in the general sense, especially in IP/WDM networks where packages are usually routed in such topologies. While grooming of traffic in such networks has been investigated recently with other heuristic algorithms,[23–26] only one paper presented a GA approach combined with clustering heuristics to deal with this problem with single- and multi-hop connections.[27] Since traffic grooming in a mesh network is much more complex than in a ring network, metaheuristic approaches are more than needed.

Due to the 'mesh' characteristic of such networks, metaheuristic approaches alone — pure GAs for instance — are not powerful enough to obtain good grooming results. An ideal way to grooming traffic in such networks might be the incorporation of metaheuristic approaches with other efficient local or greedy search mechanisms. The heuristic algorithms that have been developed so far, in Refs. 23–26 for example, could be used for this purpose. This incorporation of metaheuristics with other search heuristics is significant not only in traffic grooming but also in other optical network design problems, such as routing and wavelength assignment, protection, optimisation of wavelength converters and optical amplifiers, etc.

5.3. *Grooming of dynamic traffic*

At present, the metaheuristic algorithms developed for the grooming of dynamic traffic are in fact offline algorithms. These algorithms usually work on a set of known traffic matrices (patterns) and assign these matrices one by one to wavelengths to realise 'dynamic grooming'. Obviously, this is not dynamic grooming in its true sense. A real dynamic grooming algorithm must be able to groom dynamically to satisfy random connection requests on the network. This requires a fast online algorithm. Metaheuristic approaches are in general not suitable for this purpose since they usually need time-consuming iterations to complete a single computation. Therefore, an alternative way must be utilised to incorporate metaheuristic approaches into such a purpose.

A possible solution to this problem may be the combination of an offline static grooming algorithm with an online dynamic grooming one. The offline algorithm is used to perform a strictly nonblocking or rearrangeably nonblocking grooming for a set of estimated traffic patterns, which can be made by taking advantage of available information such as customer prescriptions, traffic projections, and historical data, etc. Based on the offline grooming results, a fast online algorithm can be used to dynamically assign traffic requests to the pre-designated wavelengths. If the current traffic cannot be assigned to these wavelengths, the algorithm can first assign a split part of the traffic to one of them, and then assign the remaining part to another wavelength with free capacity. In this way, the virtues of both approaches, metaheuristics for searching optimal grooming results and online heuristics for quickly assigning traffic to wavelengths, can be perfectly combined together.

The effectiveness of this hybrid approach can be justified by the recent observations of Elwalid *et al.*[28] In their study of routing and protection in GMPLS networks, it was found that using optimised paths computed offline to guide online path setups outperformed mere online algorithms under a wide range of operating conditions. This method was also robust to inaccuracies in the estimation of the traffic matrix. In our previous studies we have also found that offline optimisation could make full use of the dynamics of the changing traffic and make some compensation among the traffic demands between different pairs of nodes.[6,17] That is to say, it is often that at one moment some traffic requests are at their higher level with others at their lower level. At another moment this situation may be reversed. By utilising offline optimisation the traffic requests that are at their high level can be compensated by those that are at their low level, which will guarantee that the changing traffic is well fitted into the pre-designated wavelengths.

5.4. *Traffic grooming with QoS considerations*

Quality of Service (QoS) refers to the capability of a network to provide better service to selected network traffic flows with various kinds of technologies. ITU-T recommendation defines QoS as 'the collective effect of service performance which determines the degree of satisfaction of a user of the service'. The primary goal of QoS is to provide differentiated services to different priority traffic requests, and the need for internet service providers to be able to provide QoS to different customers is becoming more and more urgent.

Although QoS provisioning in the internet using IP has been a topic of active research in the last few years, QoS issues have not been considered in traffic grooming so far. The inclusion of QoS requirements into the traffic grooming process will be an interesting issue which will either reduce the usuage of resources under a certain blocking probability or lower the blocking probability under the same resource consumption. One possible way to guarantee QoS in the traffic grooming process is to assign high-priority traffic flows onto separate wavelengths and then assign other lower-priority traffic flows to share the same wavelength with it. In this way, we can preserve enough bandwidth for the high-priority traffic flows during dynamic provisioning of traffic to guarantee their low blocking probability. Some other ways of realising QoS during grooming may also be possible but researches have to be done to make the best use of a wavelength's capacity and at the same time guarantee different QoSs to different traffic flows.

6. Conclusion

We have made an extensive and comprehensive review on various metaheuristic approaches, particularly GAs, TS, and SA, to the grooming of static and dynamic traffic in WDM optical networks. Some future challenges and research directions have also been discussed in this paper. Because of the NP-hard property of grooming problems in WDM optical networks, metaheuristic approaches would be the most promising candidates for them. However so far only limited work has been done in applying these approaches to the grooming problems. On the other hand, due to the widespread use of WDM or dense WDM technologies, more and more wavelengths will be carried by a single fibre. Hence, the design of such networks will become even more difficult. With the use of generalised multi-protocol label switching (GMPLS) in optical networks, a common control plane (signalling and routing) can be provided for devices that switch in different domains: packet (IP traffic), time-division multiplexing (TDM), wavelength, waveband, and fibre.[29,30] This gives an important integrated multi-granularity routing and grooming possibility to such a system and it 'is being exploited to move deployed systems toward a flexible, "intelligent", and more autonomous multiservice optical layer, or perhaps — as Internet traffic continues to drive network design and deployment — Internet protocol (IP)/optical networks'.[31] We believe that the use of metaheuristic approaches to the grooming problems in such network systems will bring about the most cost-effective designing results for them.

References

1. A. L. Chiu and E. Modiano, Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks, *IEEE/OSA J. Lightwave Technol.* **19** (2001) 2146–2168.
2. X. J. Zhang and C. M. Qiao, An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings, *IEEE/ACM Trans. Netw.* **8** (2000) 608–617.
3. J. Wang, W. H. Cho, V. R. Vemuri and B. Mukherjee, Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections, *J. Lightwave Technol.* **19** (2001) 1645–1653.
4. R. Dutta and G. N. Rouskas, Traffic grooming in WDM networks: Past and future, *IEEE Netw.* **16**(6) (2002) 46–56.
5. R. Berry and E. H. Modiano, Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic, *IEEE J. Selected Areas Commun.* **18** (2000) 1961–1971.
6. Y. Xu, C. S. Xu and B. X. Wu, Strictly nonblocking grooming of dynamic traffic in unidirectional SONET/WDM rings using genetic algorithm, *Comput. Netw.* **41** (2003) 227–245.
7. C. M. Qiao, D. Datta, G. Ellinas, A. Gladisch and E. Modiano, Guest editorial: Special issue on WDM-based network architectures, *IEEE J. Selected Areas Commun.* **20** (2002) 1–3.
8. http://www.trafficgrooming.org/.
9. E. Modiano and P. J. Lin, Traffic grooming in WDM networks, *IEEE Commun. Mag.* **39**(7) (2001) 124–129.
10. O. Gerstel, R. Ramaswami and G. H. Sasaki, Cost effective traffic grooming in WDM rings, *IEEE INFOCOM*, San Francisco, USA (1998), pp. 69–77.
11. A. L. Chiu and E. H. Modiano, Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming, *Globecom '98*, Sydney, Australia (November 1998).
12. X. J. Zhang and C. M. Qiao, An effective and comprehensive approach to traffic grooming and wavelength assignment in SONT/WDM rings, *SPIE Proc. Conf. All-Opt. Netw.* **3531**, Boston, MA (September 1998).
13. H. W. Tzeng, J. L. Chen and N. K. Chen, Traffic grooming in WDM networks using genetic algorithm, *Proc. IEEE SMC '99 Conf.* **1**, Tokyo, Japan (October 1999), pp. 1003–1006.
14. J. Xu and Q. J. Zeng, Reducing electronic multiplexing in WDM rings: An evolutionary approach, *Proc. Int. Conf. Info-tech and Info-net* **2**, Beijing (October 2001), pp. 168–173.
15. Y. Xu, S. C. Xu and B. X. Wu, Traffic grooming in unidirectional WDM ring networks using genetic algorithms, *Comput. Commun.* **25** (2002) 1185–1194.
16. Y. Xu and S. C. Xu, Heuristic operators, redundant mapping and other issues in genetic algorithms, *Proc. CEC'01*, Seoul, Korea (May 2001), pp. 1398–1405.
17. K. H. Liu and Y. Xu, A new approach to improving the grooming performance with dynamic traffic in SONET rings, *Comput. Netw.* **46** (2004) 181–195.

18. K. H. Liu and Y. Xu, Combining GA with splitting methods to the rearrangeably nonblocking grooming of dynamic traffic in SONET ring networks, *Photonic Netw. Commun.* **10** (2005) 179–192.
19. L. Li, S. Scott and J. Deogun, Cost-effective approaches for circuit construction in WDM SONET rings, in *Proc. IASTED Int. Conf. Wireless Opt. Commun.*, Banff, Canada (July 2002), pp. 333–338.
20. O. Gerstel, P. Lin and G. Sasaki, Wavelength assignment in a WDM ring to minimize cost of embedded SONET rings, in *Proc. IEEE INFOCOM'98*, San Francisco, USA (March 1998), pp. 94–101.
21. J. Xu and Q. J. Zeng, Dynamic traffic grooming in interconnected WDM SDH/SONET rings, *Proc. SPIE* **4527** (2001) 177–186.
22. S. Zhang and B. Ramamurthy, Dynamic traffic grooming algorithms for reconfigurable SONET over WDM networks, *IEEE J. Selected Areas Commun.* **21** (2003) 1165–1172.
23. Z. Ding and H. Mounir, Clustering techniques for traffic grooming in optical WDM mesh networks, in *Proc. IEEE GLOBECOM '02* **3**, Taipei, Taiwan (November 2002), pp. 2711–2715.
24. K. Zhu and B. Mukherjee, Traffic grooming in an optical WDM mesh network, *IEEE JSAC* **20** (2002) 122–33.
25. H. Wen, R. He, L. Li and S. Wang, A plane-construction traffic grooming algorithm in WDM mesh networks, in *Proc. PDCAT'2003*, IEEE, Chengdu, P.R. China (August 2003), pp. 263–267.
26. N. Srinivas and C. S. R. Murthy, Design and dimensioning of a WDM mesh network to groom dynamically varying traffic, *Photonic Netw. Commun.* **7** (2004) 19–191.
27. C. W. Lee and E. K. Park, A genetic algorithm for traffic grooming in all-optical mesh networks, in *Proc. IEEE SMC'02* **7**, Hammamet, Tunisia (October 2002), pp. 6–9.
28. A. Elwalid, D. Mitra, I. Saniee and I. Widjaja, Routing and protection in GMPLS networks: From shortest paths to optimized designs, *IEEE/OSA J. Lightwave Technol.* **21** (2003) 2828–2837.
29. A. Banerjee, J. Drake, J. P. Lang, B. Turner, K. Kompella, and Y. Rekhter, Generalized multiprotocol label switching: An overview of routing and management enhancements, *IEEE Commun. Mag.* **39**(1) (2001) 144–150.
30. A. Banerjee, J. Drake, J. P. Lang, B. Turner, D. Awduche, L. Berger, K. Kompella and Y. Rekhter, Generalized multiprotocol label switching: An overview of signaling enhancements and recovery techniques, *IEEE Commun. Mag.* **39**(7) (2001) 144–151.
31. M. Jinno, F. J. Leonberger, M. Medard, N. Ransom and A. V. Lehmen, Guest editorial: Optical networks, *IEEE J. Lightwave Technol.* **21** (2003) 2452–2454.